arXiv:1411.6698v2 [math.PR] 14 Sep 2016

PROBABILISTIC DIVIDE-AND-CONQUER: DETERMINISTIC SECOND HALF

STEPHEN DESALVO

Abstract. We present a probabilistic divide-and-conquer (PDC) method for *exact* sampling of conditional distributions of the form $\mathcal{L}(\mathbf{X} \,|\, \mathbf{X} \in E)$, where $\mathbf{X}$ is a random variable on $\mathcal{X}$, a complete, separable metric space, and event E with $\mathbb{P}(E) \geq 0$ is assumed to have sufficient regularity such that the conditional distribution exists and is unique up to almost sure equivalence. The PDC approach is to define a decomposition of $\mathcal{X}$ via sets $\mathcal{A}$ and $\mathcal{B}$ such that $\mathcal{X} = \mathcal{A} \times \mathcal{B}$, and sample from each separately. The deterministic second half approach is to select the sets $\mathcal{A}$ and $\mathcal{B}$ such that for each element $a \in \mathcal{A}$, there is only one element $b_a \in \mathcal{B}$ for which $(a, b_a) \in E$. We show how this simple approach provides non-trivial improvements to several conventional random sampling algorithms in combinatorics, and we demonstrate its versatility with applications to sampling from sufficiently regular conditional distributions.

1. Introduction

The random sampling problem which motivates us is the following. For any integer $n \geq 1$, let $\mathbf{X}_n = (X_1, X_2, \ldots, X_n)$ denote the $\mathbb{R}^n$–valued joint distribution with coordinates independent random variables with known marginal distributions $\mathcal{L}(X_1)$, $\mathcal{L}(X_2)$, ..., $\mathcal{L}(X_n)$; denote by $\mathcal{F}$ the Borel sigma-algebra of measurable sets on $\mathbb{R}^n$. Given an event $E_n \in \mathcal{F}$, we define the distribution of $\mathbf{X}'_n$ as

$$\mathcal{L}(\mathbf{X}'_n) := \mathcal{L}\left((X_1, X_2, \ldots, X_n) \;\middle|\; \mathbf{X}_n \in E_n\right). \tag{1}$$

Many random sampling problems of interest can be described in this manner; see for example [6, 7, 24, 33].

It shall be convenient to consider in these first few sections a more abstract version of the sampling problem, although we stress that one loses very little intuition by continuing to think in terms of (1): let $\mathcal{X}$ denote a metric space, $\mathcal{F}$ the Borel sigma-algebra of measurable sets on $\mathcal{X}$, and $\mathbf{X}$ an $\mathcal{F}$-measurable random variable which induces a Radon (tight) probability measure λ on $\mathcal{X}$. Let $\mathcal{Y}$ denote another metric space, with $\mathcal{G}$ the Borel sigma-algebra of measurable sets on $\mathcal{Y}$. Suppose T is a measurable map from $(\mathcal{X}, \mathcal{F})$ into $(\mathcal{Y}, \mathcal{G})$. Then, given the event $E = \{T = t\}$, we wish to sample from

$$\mathcal{L}(\mathbf{X}') = \mathcal{L}(\,\mathbf{X} \,|\, T = t\,). \tag{2}$$

In Section 6, we give an example where this generality allows us to sample from a distribution which is not otherwise accessible by elementary conditioning.

Next, we distinguish between *approximate* and *exact* sampling. By approximate sampling, we mean an algorithm such as a (forward) Markov chain (but *not* coupling from the past [49]) or a Gibbs sampler [34], which (unless one already starts with a random sample from the

Date: September 15, 2016.

desired distribution) approaches the desired distribution "in the limit," so that after any deterministically fixed, finite time there is an inherent amount of bias in the algorithm, albeit often acceptably small for many applications. By exact sampling, we mean an algorithm which terminates in finite time with probability 1 with a sample from the given distribution. See the recent book [38] for further examples of exact sampling, where it is also referred to as *perfect simulation*.

Goal: *Generate* exact *samples as "efficiently" as possible from* $\mathcal{L}(\mathbf{X}')$.

Our view of "efficient" is the same as that specified in the introduction in [24]. That is, we address the following practical problem: how much memory and time are required to sample from $\mathcal{L}(\mathbf{X}')$? We will assume that the dominant cost of an algorithm is the number of calls to a random number generator, which produces i.i.d. uniform random values in the interval $(0, 1)$, each with a cost which is $O(1)$; we also assume that arithmetic is largely negligible, and that precision is fixed. Finally, the runtimes of our algorithms will themselves be random variables, most with unbounded ranges, and so we only consider those algorithms which terminate in finite time with probability 1. As a result, we use the *expected* runtime of an algorithm for the measure of efficiency.

When event E has positive probability, the conditional distribution is well-defined and one can aways sample from the distribution $\mathcal{L}(\mathbf{X})$ repeatedly until a variate generated lies in the set E, although the number of times to restart the algorithm may be impractical. We refer to this type of algorithm as *hard rejection sampling*, since each outcome is rejected with probability 0 or 1. A more general approach is to sample from a related distribution $\mathcal{L}(\mathbf{Y})$, and reject each outcome with a probability in $[0, 1]$, which depends on $\mathcal{L}(\mathbf{Y})$ and the observed variate; we refer to this type of algorithm as *soft rejection sampling*, since it requires some auxiliary randomness and does not always outright reject certain outcomes.

In the case when event E has probability 0, we are not aware of the existence of any analogous hard rejection sampling algorithms which can be applied generally. We are aware of several methods tailored for particular applications, for example sampling from the surface of an n–sphere, see for example [8, 21], or convex polytope sampling, see for example [20, 30], but nothing as general as hard rejection sampling which applies as readily and ubiquitously.

At this point we also remind the reader that we are not considering algorithms which generate samples from an approximate distribution in finite time, such as standard (forward) Markov chains, or the Boltzmann sampler, which considers samples of $\mathcal{L}(\mathbf{X})$ as acceptable surrogates for $\mathcal{L}(\mathbf{X}')$. Also included in this category are sampling algorithms which replace the event E by some event E' with positive probability, with $E \subset E'$. There is one exact sampling Markov chain approach we are aware of, coupling from the past [49], which typically requires some type of monotonicity in order to apply practically.

Probabilistic divide-and-conquer (PDC) is a method for exact sampling which divides a sample space into two parts, samples each part separately, and then pieces them back together to form an exact sample [5]. That is, we take as our sample space some set $\mathcal{X}$, and decompose it by sets $\mathcal{A}$ and $\mathcal{B}$ such that $\mathcal{X} = \mathcal{A} \times \mathcal{B}$; we refer to this decomposition as the *PDC division*. Then, for some random variable $\mathbf{X}$ on $\mathcal{X}$ and event E, we assume that $\mathbf{X}$ can be written as $\mathbf{X} = (A, B)$, with random variables $A \in \mathcal{A}$ and $B \in \mathcal{B}$ such that A and B are independent, and $\mathcal{L}(\mathbf{X} \,|\, E) = \mathcal{L}((A, B) \,|\, E)$. One then samples from $\mathcal{L}(A \,|\, E)$, say observing variate x, and then samples from $\mathcal{L}(B \,|\, E, A = x)$, say observing variate y;

the PDC Lemma [5, Lemma 2.1] affirms that (x, y) is indeed an exact sample from $\mathcal{L}(\mathbf{X} \mid E)$ when event E has positive probability.

Our purpose in this paper is two-fold. First, we prove in Lemma 3.1 a generalization to the PDC Lemma [5, Lemma 2.1], which contains sufficient conditions for when PDC can be applied even when event E has probability 0. The second is to expand on the utility of a particular PDC parameterization, originally exploited in [5, Section 3.3] for the random sampling of integer partitions, dubbed *deterministic second half*, which can (and should) be exploited almost as ubiquitously as hard rejection sampling. The name comes from the fact that we shall assume that sets $\mathcal{A}$ and $\mathcal{B}$ in the PDC division are chosen so that for each $a \in \mathcal{A}$, there is exactly one $b_a \in \mathcal{B}$ such that $(a, b_a) \in E$.

The outline of the paper is as follows. In Section 2, we state conditions under which the conditional distribution (2) exists and is unique up to almost sure equivalence of probability distributions. In Section 3, we prove an extended PDC lemma which allows for certain events of probability 0, and provide a general PDC algorithm for the random sampling of distributions of the form in (2), stated in terms of quantities which may not be explicitly or efficiently computable in all situations, but which does *not* yet assume a deterministic second half PDC division. In Section 4, we specialize to the deterministic second half case, and present two widely applicable PDC algorithms, one for discrete random variables and one for continuous random variables, which are stated in terms of explicitly computable quantities. In Section 5, we provide an analysis of the costs associated with the deterministic second half algorithm. In Section 6, we give an example which demonstrates that PDC is *not* simply rejection sampling, and which also highlights the subtleties involved with conditioning on events of probability 0. In Section 7, we describe several applications of PDC deterministic second half to joint distributions of continuous random variables with event E having probability 0. Finally, in Section 8, we demonstrate the effectiveness of PDC deterministic second half for the random sampling of certain combinatorial distributions.

2. Conditioning as disintegration

This section takes its name from [14], where an elegant, rigorous treatment of conditional distributions is presented: quoting from page 289, (with references for Tjur [57] and Winter [59])

> It has long bothered us (and other authors, such as Tjur, 1974 and Winter, 1979) that there should be such a wide gap between intuition and rigor in conditioning arguments. We feel that, in many statistical problems, manipulation of the conditional probability distribution is the most intuitive way to proceed. However, we mathematical statisticians are trained to treat such conditional distributions with great caution, being aware of the menagerie of nasty counterexamples – such as the Borel paradox – that warn one away from conditional distributions. ...
>
> By way of a small amount of theory and a collection of illustrative examples, in this paper we present a case that disintegrations are easy to manipulate and that they recapture some of the intuition lost by the more abstract approach, allowing guilt-free manipulation of conditional distributions.

We shall endeavor to recount the essential aspects of the framework presented in [14] so that it is apparent what conditions are sufficient for a probabilistic divide-and-conquer algorithm to be fashioned, and highly recommend a thorough reading of [14] for a more complete picture

with enlightening examples. There is also a very reasonable topological restriction on the state space, which we emphasize is likely to be satisfied in most applications, in particular those involving real-valued random variables.

We start with the definition of a disintegration, which is essentially a sigma-finite measure which acts like the elementary notion of a conditional distribution. We now introduce some definitions using standard terminology. Let $(\mathcal{X}, \mathcal{F})$ and $(\mathcal{Y}, \mathcal{G})$ denote two measurable spaces (where we have adopted standard nomenclature from [32]), and suppose T is a measurable map from $(\mathcal{X}, \mathcal{F})$ into $(\mathcal{Y}, \mathcal{G})$. Suppose λ is a probability measure on $(\mathcal{X}, \mathcal{F})$. Borrowing the notation from [14], we write $\lambda f = \int f(x)\lambda(dx)$, and for events $A \subset \mathcal{F}$ we let $\lambda(fA) = \int f(x)\mathbb{1}\{x \in A\}\lambda(dx)$. Similarly, we denote the image of measure λ by T as $T\lambda$, with $(T\lambda)B = \lambda(T^{-1}A)$ denoting the probability of event $B \subset \mathcal{B}$ whenever $TA = B$.

Next, we recall the essential characteristics of a conditional probability measure. First, we assume $\mathbb{P}$ is a probability measure on $(\mathcal{X}, \mathcal{F})$. We next assume that T only takes values in some finite subset $\mathcal{R} \subset \mathcal{Y}$ such that $\mathbb{P}\{T = t\} > 0$ for all $t \in \mathcal{R}$. Then the measure $\mathbb{P}(\cdot \,|\, T = t)$ satisfies

(a) $\mathbb{P}(A \,|\, T = t) = \dfrac{\mathbb{P}(A \cap \{T = t\})}{\mathbb{P}\{T = t\}}$, for $A \in \mathcal{F}$ and $t \in \mathcal{R}$;

(b) $\mathbb{P}(\cdot \,|\, T = t)$ is a probability measure on $(\mathcal{X}, \mathcal{F})$ for all $t \in \mathcal{R}$;

(c) $\mathbb{P}(A) = \sum_{t \in \mathcal{R}} \mathbb{P}\{T = t\}\mathbb{P}(A \,|\, T = t)$, for $A \in \mathcal{F}$.

Next is the definition of a disintegration. Let λ be a sigma-finite measure on $\mathcal{F}$ and μ be a sigma-finite measure on $\mathcal{G}$.

Definition 2.1 ([14])**.** *We say that λ has a disintegration $\{\lambda_t\}$ with respect to T and μ, or a (T, μ)-disintegration, if, for each nonnegative measurable function f on $\mathcal{X}$:*

(i) *λ_t is a probability measure on $\mathcal{F}$ concentrated on $\{T = t\}$; that is, $\lambda_t\{T \neq t\} = 0$ for μ-almost all t;*

(ii) *$t \mapsto \lambda_t(f)$ is measurable;*

(iii) *$\lambda(f) = \int_t \lambda_t(f)d\mu(t)$.*

We shall see shortly that whenever the disintegration $\{\lambda_t\}$ consists of probability measures then we have $\mu = T\lambda$. Next, let us briefly explain sufficient conditions under which such a disintegration exists and is unique; see Theorem 2.1 below. It is sufficient that the set $\mathcal{G}$ is countably generated and contains all singleton sets, and that λ is a Radon measure on a metric space. A *Radon measure* is a Borel measure for which $\lambda(K) < \infty$ for each compact K, and for any Borel set B we define $\lambda(B) = \sup_{K \subseteq B} \lambda(K)$, where the supremum is over compact sets K. As pointed out in [14], by [10, Theorem 1.4], a finite Borel measure on a complete, separable metric space is Radon (also known as *tight*). Thus, we can specialize to many common probability measures on $\mathbb{R}^n$ without difficulty.

Theorem 2.1 ([14, Existence Theorem])**.** *Let λ be a sigma-finite Radon measure on a metric space $\mathcal{X}$ and let T be a measurable map from $(\mathcal{X}, \mathcal{F})$ into $(\mathcal{Y}, \mathcal{G})$. Let μ be a sigma-finite*

measure on $\mathcal{G}$ that dominates the image measure $T\lambda$. If $\mathcal{G}$ is countably generated and contains all the singleton sets $\{t\}$, then λ has a (T,μ)-disintegration. The λ_t measures are uniquely determined up to an almost sure equivalence: if $\{\lambda_t^*\}$ is another (T,μ)-disintegration then $\mu\{t \in \mathcal{Y} : \lambda_t \neq \lambda_t^*\} = 0$.

The next two theorems now complete the picture by providing the conditions for which the disintegration consists of probability measures, and also when the densities exist.

Theorem 2.2 ([14])**.** *Let λ have a (T,μ)-disintegration $\{\lambda_t\}$, with λ and μ each sigma-finite.*

(i) *The image measure $T\lambda$ is absolutely continuous with respect to μ, with density $\lambda_t\mathcal{X}$.*
(ii) *The measures $\{\lambda_t\}$ are finite for μ-almost all t if and only if $T\lambda$ is sigma-finite.*
(iii) *The measures $\{\lambda_t\}$ are probabilities for μ-almost all t if and only if $\mu = T\lambda$.*
(iv) *If $T\lambda$ is sigma-finite then $(T\lambda)\{\lambda_t\mathcal{X} = 0\} = 0$ and $(T\lambda)\{\lambda_t\mathcal{X} = \infty\} = 0$. For $T\lambda$-almost all t, the measures*

$$\tilde{\lambda}_t(\cdot) = \frac{\lambda_t(\cdot)}{\lambda_t\mathcal{X}}\,\{0 < \lambda_t\mathcal{X} < \infty\}$$

are probabilities that give a $(T, T\lambda)$-disintegration of λ.

Theorem 2.3 ([14])**.** *Let λ have a (T,μ)-disintegration $\{\lambda_t\}$, and let ρ be absolutely continuous with respect to λ with a finite density $r(x)$, with each of λ, μ, and ρ sigma-finite.*

(i) *The measure ρ has a (T,μ)-disintegration $\{\rho_t\}$ where each ρ_t is dominated by the corresponding λ_t, with density $r(x)$.*
(ii) *The image measure $T\rho$ is absolutely continuous with respect to μ, with density $\lambda_t r$.*
(iii) *The measures $\{\rho_t\}$ are finite for μ-almost all t if and only if $T\rho$ is sigma-finite.*
(iv) *The measures $\{\rho_t\}$ are probabilities for μ-almost all t if and only if $\mu = T\rho$.*
(v) *If $T\rho$ is sigma-finite then $(T\rho)\{\lambda_t r = 0\} = 0$ and $(T\rho)\{\lambda_t r = \infty\} = 0$. For $T\rho$-almost all t, the measures defined by*

$$\tilde{\rho}_t(f) = \frac{\lambda_t(f\,r)}{\lambda_t(r)}\,\{0 < \lambda_t(r) < \infty\} \tag{3}$$

are probabilities that give a $(T, T\rho)$-disintegration of ρ.

Again quoting from [14]:

> The simple formula (3) is the general version of the familiar method for calculating conditional densities as a ratio of joint density to marginal density. It is more useful than the familiar formula because it does not require the conditioning variable to be a coordinate projection on a Euclidean space with Lebesgue measure playing the role of λ.

3. A PDC lemma with disintegrations

We start with a sample space $\mathcal{X}$ with Borel sigma-algebra $\mathcal{F}$; random variable $\mathbf{X} \in \mathcal{F}$ which induces a probability measure $\mathbb{P}$ on $\mathcal{X}$; $\mathcal{F}$-measurable event E; and we wish to sample from the distribution in (2). We write the sample space as $\mathcal{X} = \mathcal{A} \times \mathcal{B}$, with random variables $A \in \mathcal{A}$ and $B \in \mathcal{B}$ defined such that they are independent, and (2) can be written as $\mathcal{L}((A,B) \,|\, (A,B) \in E)$. The most generic approach to sampling from the distribution (2) is hard rejection sampling, given by Algorithm 1. The generic PDC random sampling algorithm is given in Algorithm 2. Note that, when $\mathbb{P}(\mathbf{X} \in E) = 0$, hard rejection sampling does not terminate in finite time with probability 1.

Algorithm 1 Hard rejection sampling of $\mathcal{L}(\,(A,B)\,|\,(A,B)\in E\,)$

1. Generate sample from $\mathcal{L}(A)$, call it a.
2. Generate sample from $\mathcal{L}(B)$, call it b.
3. Check if $(a,b)\in E$; if so, return (a,b), otherwise restart.

Algorithm 2 Probabilistic Divide-and-Conquer sampling of $\mathcal{L}(\,(A,B)\,|\,(A,B)\in E\,)$

1. Generate sample from $\mathcal{L}(A\,|\,E)$, call it x.
2. Generate sample from $\mathcal{L}(B\,|\,E, A=x)$ call it y.
3. Return (x,y).

The PDC Lemma [5, Lemma 2.2] affirms that the resulting pair (x,y) from Algorithm 2 is an exact sample from $\mathcal{L}(\mathbf{X}')$ when $\mathbb{P}(\mathbf{X}\in E)>0$. We now generalize this lemma to include certain events for which $\mathbb{P}(\mathbf{X}\in E)=0$. Recall the notation from the previous section: Let $(\mathcal{X},\mathcal{F})$ and $(\mathcal{Y},\mathcal{G})$ denote two measurable spaces, and suppose T is a measurable map from $(\mathcal{X},\mathcal{F})$ into $(\mathcal{Y},\mathcal{G})$.

Lemma 3.1. *Suppose*

(i) $\mathbb{P}$ *is a given probability Radon measure on* $(\mathcal{X},\mathcal{F})$*;*
(ii) *the image measure* $T\,\mathbb{P}$ *is a probability measure on* $(\mathcal{Y},\mathcal{G})$*, where* $\mathcal{G}$ *is countably generated and contains all the singleton sets* $\{t\}$*.*

Then $\mathbb{P}$ *has disintegration probability measure* $\mathbb{P}_t := \mathbb{P}(\,\cdot\,|\,T=t)$*, which is unique up to almost sure equivalence of probability measures, and for events* E *of the form* $E=\{T=t\}$*, we have* $\mathcal{L}(\boldsymbol{X}')$ *is well-defined and unique on events of positive* $\mathbb{P}_t$ *measure.*

Furthermore, assume

(a) *there are metric spaces* $\mathcal{A}$ *and* $\mathcal{B}$ *such that* $\mathcal{X}=\mathcal{A}\times\mathcal{B}$*;*
(b) *there are probability measures* $\mathbb{P}_{\mathcal{A}} := \pi_{\mathcal{A}}\mathbb{P}$ *and* $\mathbb{P}_{\mathcal{B}} := \pi_{\mathcal{B}}\mathbb{P}$*, such that* $\mathbb{P}=\mathbb{P}_{\mathcal{A}}\times\mathbb{P}_{\mathcal{B}}$*, corresponding to measurable random variables* $A\in\mathcal{A}$ *and* $B\in\mathcal{B}$ *which are independent with distributions also denoted by* $\mathcal{L}(A)$ *and* $\mathcal{L}(B)$*, respectively;*
(c) *for each* $a\in\mathcal{A}$*, the map* $T_a(\cdot) := T(a,\cdot)$ *is measurable (i.e.,* T_a *is the map* T *holding the coordinate* $\mathcal{A}$ *fixed), and we let* t_a *denote the set of values such that* $\{T=t, A=a\}=\{T_a\in t_a\}$ *holds.*

Then the following distributions exist and are unique up to almost sure equivalence:

$$\mathcal{L}(X) := \mathcal{L}(A\,|\,T=t), \qquad \mathcal{L}(Y\,|\,X=x) := \mathcal{L}(B\,|\,T=t, A=a) = \mathcal{L}(B\,|\,T_a(B)\in t_a);$$

Finally, we have $\mathcal{L}(X,Y)=\mathcal{L}(\boldsymbol{X}')$*.*

Proof. First, we demonstrate that all distributions exist and are unique up to almost sure equivalence. That $\mathbb{P}$ has disintegration probability measure $\mathbb{P}_t$ is a consequence of Theorem 2.2, which also implies that $\mathcal{L}(\mathbf{X}')$ is well-defined and unique up to almost sure equivalence. By projecting $\mathbf{X}$ onto $\mathcal{A}$, we conclude the same holds for $\mathcal{L}(X)$ (see for example [14, Example 2]), and denote this measure by $\mathbb{P}_X$. By the same argument, $\mathcal{L}(B\,|\,T=t)$ also exists and is unique up to almost sure equivalence. Then, under the assumption that T_a is measurable for each $a\in\mathcal{A}$, there exists a disintegration probability measure for $\mathcal{L}(B\,|\,T=t)$ under map T_a, i.e., $\mathcal{L}(B\,|\,T=t, A=x)$ for each $x\in A$, which is precisely $\mathcal{L}(Y\,|\,X=x)$; we denote this measure by $\mathbb{P}_{Y,x}$.

Next, we show $\mathcal{L}(X,Y) = \mathcal{L}(\mathbf{X}')$. For all nonnegative measurable $f : \mathcal{A} \times \mathcal{B} \to \mathbb{R}$, by property (iii) of Definition 2.1, applied once to $\mathbb{P}$ and a second time to $\mathbb{P}_t$ (see also [14, Example 4]), we have

$$\begin{aligned}\int f(A,B)\,(\mathbb{P}_{\mathcal{A}} \times \mathbb{P}_{\mathcal{B}})(dA \times dB) &= \int_t \left(\int f(A,B)\ \mathbb{P}_t(dA \times dB) \right) (T\mathbb{P})(dt) = \\ &= \int_t \left(\int \left(\int f(A,B)\ \mathbb{P}_{Y,x}(dB) \right) \mathbb{P}_X(dA) \right) (T\mathbb{P})(dt) = \\ &= \int_t \left(\int f(X,Y)\ (\mathbb{P}_X \times \mathbb{P}_Y)(dX \times dY) \right) (T\mathbb{P})(dt).\end{aligned}$$

Then, for all bounded measurable $g : \mathcal{A} \times \mathcal{B} \to \mathbb{R}$, we separate g into its positive and negative parts and apply the same argument, which completes the proof. □

Before we specialize to a more concrete setting, let us first introduce and extend a particularly fruitful PDC algorithm which uses soft rejection sampling, see [5, Algoirthm 3]. Recall that soft rejection sampling samples a related distribution $\mathcal{L}(\mathbf{Y})$, and introduces an auxiliary random variable U, uniform over the interval $(0,1)$, which is used to reject samples of $\mathcal{L}(\mathbf{Y})$ so that they appear in their correct proportion in $\mathcal{L}(\mathbf{X}')$. In our application, we will apply soft rejection sampling to $\mathcal{L}(A \,|\, T = t)$ using samples from $\mathcal{L}(A)$. We define the function $q(a)$, $a \in \mathcal{A}$, as the correct proportion for using soft rejection sampling in this manner.

We then define the *rejection function* as

$$s(a) := \frac{q(a)}{\sup_{\ell \in \mathcal{A}} q(\ell)}, \qquad a \in \mathcal{A}. \tag{4}$$

(In the event that $q(a) = 0$ for all $a \in \mathcal{A}$, we assign $s(a) = 0$.) In other words, we will sample from the distribution $\mathcal{L}((A,U) \,|\, U < s(a))$, where $s : \mathcal{A} \to [0,1]$ is a measurable function, and with s chosen as in Equation (4), the first coordinate has distribution $\mathcal{L}(A \,|\, T = t)$. Note, however, that we need some additional assumptions so that Algorithm 3 below terminates in finite time with probability 1:

$$0 < \sup_{\ell \in \mathcal{A}} q(\ell) < \infty; \tag{A1}$$

$$\mathbb{P}_{\mathcal{A}}\{a \in \mathcal{A} : q(a) > 0\} > 0. \tag{A2}$$

Algorithm 3 PDC sampling from $\mathcal{L}((A,B) \,|\, (A,B) \in E)$ using soft rejection sampling

0. Assume the conditions of Lemma 3.1, and in addition assume (A1) and (A2).
1. Generate sample from $\mathcal{L}(A)$, call it a.
2. Reject a with probability $1 - s(a)$, where $s(a)$ is given in (4); otherwise, restart.
3. Generate sample from $\mathcal{L}(B \,|\, (a,B) \in E)$, call it y.
4. Return (a,y).

Remark 3.1. In the event that $\sup_{\ell \in \mathcal{A}} q(\ell)$ is not practical to compute efficiently, we may replace this quantity in Equation (4) with *any* upper bound C satisfying $\sup_{\ell \in \mathcal{A}} q(\ell) \le C < \infty$, and Algorithm 3 is still an exact sampling algorithm. Of course, as is well known (see for example [58]), the efficiency of the algorithm diminishes the larger C is taken, in the

sense that the expected number of rejections before we accept a sample is multiplied by the corresponding quotient $C/\sup_{\ell\in\mathcal{A}} q(\ell)$, and so it is optimal to take $C = \sup_{\ell\in\mathcal{A}} q(\ell)$.

Of course, there is still the matter of how to sample from $\mathcal{L}(B \,|\, (a, B) \in E)$ for each $a \in \mathcal{A}$. An often optimal divide-and-conquer strategy fashions sets $\mathcal{A}$ and $\mathcal{B}$ so that $\mathcal{L}(B \,|\, (a, B) \in E)$ is a smaller version of the original sampling problem $\mathcal{L}(\mathbf{X} \,|\, \mathbf{X} \in E)$, split roughly in half, and recursively repeats until a simple base case is reached; this is referred to as self-similar, recursive PDC in [5, Section 3.5], and it was used to yield an asymptotically efficient sampling algorithm for the random sampling of integer partitions. The fashioning of a self-similar PDC algorithm, however, requires the efficient computation of the rejection function $s(a)$, which can be computed efficiently in the case of integer partitions, see Remark 8.2; in general, however, the efficient computation of $s(a)$ is not a trivial matter, which is why we specialize to the more practical deterministic second half setting.

4. PDC with deterministic second half

The efficient computability of $s(a)$ in Algorithm 3 is a vital consideration when fashioning a PDC division. It is with this consideration in mind that we revisit a PDC division first introduced in [5, Section 3.3] called deterministic second half. For each $a \in \mathcal{A}$, define the set $E_a := \{b \in \mathcal{B} : (a, B) \in E\}$ as the set of all $b \in \mathcal{B}$ which can be paired with a given $a \in \mathcal{A}$ and satisfy $(a, b) \in E$. Suppose $\mathcal{A}$ and $\mathcal{B}$ are chosen so that for each $a \in \mathcal{A}$, we have $|E_a| = 1$. This type of PDC division was shown to provide surprisingly large speedups to hard rejection sampling for many interesting examples in [5, Section 3.3], while keeping $s(a)$ practical to compute. We now generalize this approach.

In addition to the assumptions in Lemma 3.1, we now add assumption (DSH), which makes the distribution $\mathcal{L}(B \,|\, (a, B) \in E)$ trivial; that is, the second stage of the algorithm is completed by a uniquely determined value.

$$\text{For each } a \in \mathcal{A}, \quad |E_a| = |\{b_a\}| = 1. \tag{DSH}$$

To reiterate, the assumption (DSH) is an abbreviation for *deterministic second half*, and is not required for PDC in general, nor for Lemma 3.1. In terms of Algorithm 3 it is provably faster than hard rejection sampling (with respect to a costing scheme which we make precise in Section 5), and in many cases of interest the rejection function is explicit and efficient to compute.

At this point we also specialize to the parameterization in (1), since it is indicative of many interesting examples. That is, we assume a sample space $\mathcal{X} = \mathbb{R}^n$, with $\mathbf{X} = (X_1, X_2, \ldots, X_n)$ a joint distribution of real-valued *independent* random variables, where each X_i has a corresponding marginal distribution $\mathcal{L}(X_i)$, sometimes written as $\mathbb{P}_{X_i}$. In addition, we assume that $T \equiv T(X_1, \ldots, X_n)$ is some measurable function which satisfies the conditions of Lemma 3.1.

Example 4.1. *Suppose $X_1, X_2, \ldots, X_n$ are discrete random variables. Let $T = \sum_{i=1}^n X_i$, so that $T\,\mathbb{P} = \mathbb{P}_{X_1} * \mathbb{P}_{X_2} * \cdots * \mathbb{P}_{X_n}$ is the convolution of measures. Then for any fixed $t \in \mathbb{R}$ in the range of T, taking $E = \{T = t\}$, we have that the distribution* (1) *is well-defined by Theorem 2.1. Since t is uniquely determined given any $n-1$ of the variables $X_1, \ldots, X_n$, a valid PDC deterministic second half approach would be to select an index i, say, e.g., $i = 1$, and let $\mathcal{A} = \mathbb{R}^{n-1}$, $\mathcal{B} = \mathbb{R}$, with random variables $A = (X_2, \ldots, X_n)$ and $B = X_1$. Applying*

Algorithm 3, the rejection function is given by

$$s((y_2,\ldots,y_n)) = \frac{\mathbb{P}\left(X_1 = t - \sum_{i=2}^n y_i\right)}{\sup_\ell \mathbb{P}(X_1=\ell)}, \qquad (y_2,\ldots,y_n)\in\mathcal{A}.$$

Example 4.2. *Suppose in Example 4.1 we instead take $X_1, X_2, \ldots, X_n$ to be continuous random variables with a density, say we denote f_{X_1} for the density of X_1. Applying Algorithm 3, the rejection function is then given by*

$$s((y_2,\ldots,y_n)) = \frac{f_{X_1}(t - \sum_{i=2}^n y_i)}{\sup_\ell f_{X_1}(\ell)}, \qquad (y_2,\ldots,y_n)\in\mathcal{A}.$$

Examples 4.1 and 4.2 form the foundation of our intuition, but the form of T is needlessly specific, and so we conclude the theoretical treatment of PDC deterministic second half with a generalization below that we have found both concrete and adequate for almost all applications encountered. Even though, as mentioned in the quotation at the end of Section 2, we are not restricted to coordinate projections and Euclidean spaces, many natural PDC divisions are in fact coordinate projections on $\mathbb{R}^k$, and so we have endeavored to make this final treatment the most widely accessible and applicable.

To simplify notation, when the values of the parameters are understood from context we adopt the conventions $\mathbf{X}' \equiv \mathbf{X}'_n$, $E \equiv E_n$, $\mathbf{X} \equiv \mathbf{X}_n \equiv (X_1,\ldots,X_n)$, $[n] = \{1,\ldots,n\}$. In order to describe the class of divisions, we now state the following definitions.

Definition 4.1. *Let $I = \{i_1, i_2, \ldots\} \subset [n]$ denote some subset of indices, with $i := |I|$.*

(1) Let $X_I := \pi_I(\boldsymbol{X}) = (X_i)_{i\in I}$ denote the $\mathbb{R}^i$–valued projection;

(2) let $X^{(I)} := \pi_{[n]\setminus I}(\boldsymbol{X})$ denote the $\mathbb{R}^{n-i}$–valued projection;

(3) let $\sigma_I : \mathbb{R}^{n-i}\times\mathbb{R}^i \to \mathbb{R}^n$ denote the (unique) rearrangement of the combined set of elements in x and y in the following manner: for $x = (x_1,\ldots,x_{n-i}) \in \mathbb{R}^{n-i}$ and $y = (y_1,\ldots,y_i \in \mathbb{R}^i$, $z = \sigma_I(x,y)$ is the vector for which $z_{i_j} = y_j$ for $j = 1,\ldots,i$, and which keeps the original order of $x_1,\ldots,x_{n-i}$ in z, i.e., for all indices $1 \le j_1 < j_2 \le n-i$ for which $x_{j_1} = z_{\ell_1}$ and $x_{j_2} = z_{\ell_2}$, we have $1 \le \ell_1 < \ell_2 \le n$. In other words, σ_I is defined so that we have $\{X\in E\} = \{\sigma_I(X^{(I)}, X_I) \in E\}$;

*(4) define the I-*completable-set *of E_n as*

$$E^{(I)} := \pi^{-1}_{[n]\setminus I}(E_n) = \{x\in\mathbb{R}^{n-i} : \exists y\in\mathbb{R}^i \text{ such that } \sigma_I(x,y)\in E_n\};$$

*(5) define the I-*section of E_n given $x^{(I)} \in E^{(I)}$ *as*

$$E_I \equiv E_{I|x^{(I)}} := \pi_I^{-1}(E_n \,|\, x^{(I)}) = \{y\in\mathbb{R}^i : \sigma_I(x^{(I)},y)\in E_n\};$$

(6) define $T_I \equiv T_I(X_I|x^{(I)}) := T|_{E_{I|x^{(I)}}}(\sigma_I(x^{(I)}, X_I))$.

To aid in the presentation of our algorithms, we also define the 'cemetery' state Δ (see, e.g., [26, Section 3.1]) to be such that $S\Delta = \Delta$ for all transformations S, with $\mathbb{P}(X=\Delta) = 0$ and the density function of X, if it exists, evaluated at Δ equals 0, for all random variables X. Under Assumption (DSH), for each $x^{(I)} \in E^{(I)}$, we define $y_I \equiv y_I(x^{(I)}) \in E_I$ to be the unique completion such that $\sigma_I(x^{(I)}, y_I) \in E$, and for each $x^{(I)} \notin E^{(I)}$, we define $y_I(x^{(I)}) = \Delta$. In addition, the random variable T_I is one-to-one on E_I, and so we define $t_I \equiv t_I(x^{(I)}) := T_I\, y_I$ as the unique image of the point y_I under map T_I. *We shall use the above definitions of y_I and t_I in all subsequent analysis and examples whenever assumption* (DSH) *is in effect.*

We now present two concrete applications of PDC with deterministic second half we believe are of most use in practical applications: one for discrete random variables and the other

for continuous random variables. See Section 6 for an application which uses Algorithm 2 directly. For continuous random variables, we now utilize Theorem 2.3, which implies the existence of the density of a disintegration under conditions which will be satisfied for our application. We will denote the density of a random variable X by f_X. In what follows, U will denote a uniform random variable between 0 and 1, independent of all other random variables, and u an observed variate from this distribution.

Algorithm 4 PDC with deterministic second half for discrete random variables

Input:
Distributions $\mathcal{L}(X_1), \ldots, \mathcal{L}(X_n)$;
measurable function T, $t \in \text{range}(T)$;
index set $I \subset \{1, \ldots, n\}$.
Output:
A sample from $\mathcal{L}((X_1, \ldots, X_n)|T = t)$.
Assumptions:
(A1), (A2), (DSH), and $X_1, \ldots, X_n, T$ are discrete random variables.

Sample from $\mathcal{L}(X^{(I)})$, denote the observation by $x^{(I)}$.
if $x^{(I)} \in E^{(I)}$ and $u < \frac{\mathbb{P}(X_I = y_I)}{\max_\ell \mathbb{P}(X_I = \ell)}$ **then**
return $\sigma_I(x^{(I)}, y_I)$
else
restart
end if

Algorithm 5 PDC with deterministic second half for continuous random variables

Input:
Distributions $\mathcal{L}(X_1), \ldots, \mathcal{L}(X_n)$;
measurable function T, $t \in \text{range}(T)$;
index set $I \subset \{1, \ldots, n\}$.
Output:
A sample from $\mathcal{L}((X_1, \ldots, X_n)|T = t)$.
Assumptions:
(A1), (A2), (DSH), and $X_1, \ldots, X_n, T$ are continuous random variables.

Sample from $\mathcal{L}(X^{(I)})$, denote the observation by $x^{(I)}$.
if $x^{(I)} \in E^{(I)}$ and $u < \frac{f_{T_I}(t_I)}{\sup_\ell f_{T_I}(\ell)}$ **then**
return $(x^{(I)}, y_I)$
else
restart
end if

Each of these algorithms follows by Lemma 3.1 and Algorithm 3. However, we shall write out explicitly the steps with which to derive the rejection function in order to draw attention to a key difference between the discrete and continuous versions; see also Remark 4.1 below.

Note that while our rejection probability in Algorithm 5 is stated in terms of the distribution $\mathcal{L}(T_I \mid X^{(I)} = x^{(I)})$, we return the value y_I, which is the corresponding value in the range of X_I.

Theorem 4.1. *Algorithm 4 samples from the distribution given in Equation* (1).

Proof. We demonstrate that $\frac{P(X_I=y_I)}{\max_\ell(P(X_I=\ell))}$ is the right proportion.

We let h denote the probability mass function of $\mathcal{L}(X^{(I)} \mid E_n)$, and g denote the probability mass function of $\mathcal{L}(X^{(I)})$. Our rejection proportion is of the form: suppose we observe state j under the distribution g, then we reject if

$$u > \frac{h(j)}{Cg(j)},$$

where C is any constant such that

$$h(\ell) \leq Cg(\ell), \qquad \text{for all states } \ell.$$

The quantity C is the expected number of iterations of the acceptance/rejection procedure before we accept a sample, see for example [20], and so in particular we would like to find the smallest C. Since our distributions are already specified, we obtain

$$\frac{1}{C} = \min_\ell \frac{g(\ell)}{h(\ell)} = \min_\ell \frac{\mathbb{P}(T = k)}{\mathbb{P}(T = k|X^{(I)} = \ell)} = \frac{\mathbb{P}(T = k)}{\max_\ell \mathbb{P}(T_I(X_I|\ell) = t_I(\ell)|X^{(I)} = \ell)},$$

and our rejection step reduces to

$$u > \frac{h(j)}{Cg(j)} = \frac{\mathbb{P}(T_I(X_I|j) = t_I(j))}{\max_\ell \mathbb{P}(T_I(X_I|\ell) = t_I(\ell))}.$$

By assumption (DSH), once we accept $x^{(I)}$, the completion y_I is unique. Since T_I is discrete, we have

$$\frac{\mathbb{P}(T_I = t_I(j))}{\max_\ell \mathbb{P}(T_I = t_I(\ell))} = \frac{\mathbb{P}(X_I = y_I(j))}{\max_\ell \mathbb{P}(X_I = \ell)}.$$

□

Theorem 4.2. *Algorithm 5 samples from the distribution given in Equation* (1).

Proof. The proof follows in a similar manner as the proof of Theorem 4.1, with probabilities replaced with probability density functions *where appropriate*, which are guaranteed to exist by Theorem 2.3 since we assume T is continuous. We have

$$\frac{1}{C} = \inf_\ell \frac{g(\ell)}{h(\ell)} = \inf_\ell \frac{f_T(k)}{f_{T_I|\ell}(t_I(\ell))} = \frac{f_T(k)}{\sup_\ell f_{T_I|\ell}(t_I(\ell))},$$

and the rejection step reduces to

$$u > \frac{h(j)}{Cg(j)} = \frac{f_{T_I|j}(t_I(j))}{\sup_\ell f_{T_I|\ell}(t_I(\ell))}.$$

□

Remark 4.1. There is an important difference between algorithms 4 and 5. When X_I and T_I are discrete, the random variables can be rearranged *before* the soft rejection step from $\{T_I = t_I(x^{(I)})\}$ to $\{X_I = y_I(x^{(I)})\}$, whereas when X_I and T_I are continuous, one would have to apply a standard change of variables formula in order to determine the rejection probability in terms of the density of X_I; see for example Section 7.5.

5. Speedup analysis

The hard rejection sampling acceptance condition is $\{\mathbf{X} \in E\}$, which in the discrete setting may be written as

$$\{X^{(I)} \in E^{(I)} \text{ and } U < P(X_I = y_I)\}.$$

The acceptance condition for Algorithm 4 is

$$\left\{X^{(I)} \in E^{(I)} \text{ and } U < \frac{\mathbb{P}(X_I = y_I)}{\max_\ell \mathbb{P}(X_I = \ell)}\right\}. \tag{5}$$

It is easy to see that there exists a coupling of the random variables such that all events giving an acceptance in the hard rejection sampling algorithm would also be accepted in the PDC deterministic second half algorithm. The added efficiency in this case comes from the use of soft rejection sampling, which enlarges the space by the factor $\max_\ell \mathbb{P}(X_I = \ell)^{-1}$.

For continuous random variables, soft rejection sampling transforms the otherwise expected infinite-time hard rejection sampling algorithm into an acceptance condition of the form

$$\left\{X^{(I)} \in E^{(I)} \text{ and } U < \frac{f_{T_I|X^{(I)}}(t_I)}{\sup_\ell f_{T_I|X^{(I)}}(\ell)}\right\}, \tag{6}$$

which is an event of *positive* probability under assumptions (A1), (A2), (DSH).

Since the memory requirements for hard rejection sampling and the algorithms presented are on the same order of magnitude, we focus solely on run–time. We assume that arithmetic operations are negligible, whereas the cost of generating a single random uniform variate from a given interval is $O(1)$, and the cost of generating n independent uniform random variables from a given interval is $O(n)$, regardless of the magnitudes of the values[1]. Implicitly, we also assume, quite critically, but also quite reasonably for many applications, that computing y_I and t_I is always $O(1)$; that is, we assume that completing the sample is indeed a trivial matter.

Definition 5.1. *For a given algorithm P, which generates a sample from a distribution $\mathcal{L}(\boldsymbol{X})$, we denote the* expected *time of completion by* $\mathrm{Time}_P(\boldsymbol{X})$. *When there is no subscript, we assume there is a default direct sampling method available.*

When P is the hard rejection sampling algorithm, we denote the expected time of completion by $\mathrm{Time}_{\mathrm{rej}}(\boldsymbol{X}')$. *When* $0 < \mathrm{Time}_P(\boldsymbol{X}) < \infty$, *the* speedup *of algorithm P relative to hard rejection sampling is defined by*

$$\text{speedup} := \frac{\mathrm{Time}_{\mathrm{rej}}(\boldsymbol{X}')}{\mathrm{Time}_P(\boldsymbol{X}')}. \tag{7}$$

Remark 5.1. We have

$$\mathrm{Time}_{\mathrm{rej}}(\mathbf{X}') = O\left(\frac{\mathrm{Time}(\mathbf{X})}{\mathbb{P}(\mathbf{X} \in E)}\right).$$

For all algorithms P with $0 < \mathrm{Time}(P) < \infty$, we have speedup $\in (0, \infty]$, with speedup > 1 representing an improvement to hard rejection sampling. In the case when hard rejection sampling does not terminate in finite time with probability 1, we write speedup $= \infty$.

[1]One might think of this as fixed floating-point precision implemented on a computer.

Theorem 5.1. *Let* $\mathrm{Time}_{\mathrm{DDSH}}(\boldsymbol{X}')$ *denote the expected run-time of Algorithm 4. We have*

$$\mathrm{Time}_{\mathrm{DDSH}}(\boldsymbol{X}') = O\left(\mathrm{Time}_{\mathrm{rej}}(\boldsymbol{X}') \max_{\ell} \mathbb{P}(X_I = \ell)\right),$$

whence

$$\text{speedup} = \Omega\left(\left(\max_{\ell} \mathbb{P}(X_I = \ell)\right)^{-1}\right). \tag{8}$$

Proof. Recall the optimal value of C is given by

$$C = \frac{\max_\ell \mathbb{P}(X_I = \ell)}{\mathbb{P}(T = k)}.$$

The cost of this algorithm with the optimal C is then

$$\begin{aligned}\mathrm{Time}\left(X^{(I)}\right) C &= O\left(\frac{\mathrm{Time}\left(X^{(I)}\right)}{P(T=k)} \max_{\ell} \mathbb{P}(X_I = \ell)\right) \\ &= O\left(\mathrm{Time}_{\mathrm{rej}}\left(\mathbf{X}'\right) \max_{\ell} \mathbb{P}(X_I = \ell)\right),\end{aligned}$$

which by Equation (7) implies the speedup is $\Omega(\max_\ell \mathbb{P}(X_I = \ell))^{-1}$. □

Remark 5.2. Theorem 5.1 indicates that the optimal choice of I in Algorithm 4 is one that minimizes the maximal point mass in the distribution of X_I.

Theorem 5.2. *Let* $\mathrm{Time}_{\mathrm{CDSH}}(\boldsymbol{X}')$ *denote the expected run-time of Algorithm 5. We have*

$$\mathrm{Time}_{\mathrm{CDSH}}(\boldsymbol{X}') = O\left(\mathrm{Time}_{\mathrm{rej}}\left(X^{(I)} | E^{(I)}\right) \sup_{\ell} f_{T_I|\ell}(t_I(\ell))\right).$$

Also, speedup $= \infty$.

Proof. The optimal value of C is given by

$$C = \frac{\sup_\ell f_{T_I|\ell}(t_I(\ell))}{f_T(k)}.$$

The cost of this algorithm with the optimal C is then

$$\begin{aligned}\mathrm{Time}\left(X^{(I)}\right) C &= O\left(\mathrm{Time}\left(X^{(I)}\right) \frac{\sup_\ell f_{T_I|\ell}(t_I(\ell))}{f_T(k)}\right) \\ &= O\left(\mathrm{Time}_{\mathrm{rej}}\left(X^{(I)} | E^{(I)}\right) \sup_{\ell} f_{T_I|\ell}(t_I(\ell))\right).\end{aligned}$$

By assumption, $E = \{T = t\}$ for some continuous random variable T with a density and $t \in \mathrm{range}(T)$, and so $\mathbb{P}(\mathbf{X} \in E) = 0$, whence, since Algorithm 5 has finite expected time, we have speedup $= \infty$. □

6. An illustrative example

This section highlights the fact that PDC is not simply rejection sampling. It also motivates the use of disintegrations from Section 2 rather than elementary conditioning.

An example which demonstrates that care must be taken when conditioning on events of probability 0 is given by the following example of [50] (see also [13, Section 4.9.3]). The problem stated on a particular exam is as follows: *If U and V are independent standard*

normals, what is the conditional distribution of V given that $V = U$? The three distinct ways in which this problem was solved started with the following joint distributions:

(1) $\big(\ (U,V) \ \big|\ U - V = 0 \ \big)$;
(2) $\big(\ (U,V) \ \big|\ \frac{U}{V} = 1 \ \big)$;
(3) $\big(\ (U,V) \ \big|\ \mathbb{1}(U = V) \ \big)$.

In our notation, this is the same as $((U,V) \,|\, T = t)$, where

(1) $T = U - V$ and $t = 0$;
(2) $T = V/U$ and $t = 1$;
(3) $T = \mathbb{1}(U = V)$ and $t = 1$.

To apply Algorithm 5, we first sample U from a standard normal distribution, and apply a rejection depending on the distribution $\mathcal{L}(T_I \,|\, X^{(I)} = a)$; we have

(1) $T_I = V - a$, reject if $u > e^{-a^2/2}$;
(2) $T_I = a/V$, reject if $u > |a|\, e^{-a^2}/\sqrt{2\,e}$;
(3) $T_I = \mathbb{1}(V = a)$, Not applicable.

Note that the last case, $T = \mathbb{1}(U = V)$, does not satisfy the assumptions of Algorithm 4 or Algorithm 5; that is, $\mathbb{P}(\mathbf{X} \in E)$ is not an event of positive probability, nor is T a random variable with a density. Rather than applying Algorithm 5, we can instead determine the conditional distribution $\mathcal{L}(V \,|\, T = t)$ directly, as the original problem demands, sample according to that distribution, and then appeal to Lemma 3.1 and Algorithm 2 directly.

(1) $T_I = V - a$, $f_{V|T=0}(v) = e^{-v^2}/\sqrt{\pi}$, $-\infty < v < \infty$;
(2) $T_I = a/V$, $f_{V|T=1}(v) = |v| e^{-v^2}$, $-\infty < v < \infty$;
(3) $T_I = \mathbb{1}(V = a)$, $f_{V|T=1}(v) = e^{-v^2/2}/\sqrt{2\pi}$, $-\infty < v < \infty$.

In each of these cases, the density $f_{V|T}$ is guaranteed to exist by Theorem 2.3, and we have $\mathcal{L}(V \,|\, T = t, U = a)$ is a point mass at a, $a \in \mathbb{R}$.

7. Applications of a theoretical nature

7.1. **Uniform weight.** This section contains a compelling application of PDC deterministic second half, one which is implicit in many other applications. That is, when the random variable T_I assigns the same weight to t_I for all $x^{(I)} \in E^{(I)}$, then each sample generated from $\mathcal{L}(X^{(I)})$ is accepted with the same proportion, which can be scaled up to 1, giving a rejection probability of 0.

Theorem 7.1. *Algorithm 6 generates a sample from the distribution $\mathcal{L}(\boldsymbol{X}')$, with expected runtime $O(\mathrm{Time}(X^{(I)})\ \mathbb{P}(X^{(I)} \in E^{(I)})^{-1})$.*

Proof. If $X^{(I)} \notin E^{(I)}$, we reject with probability 1. Assuming $X^{(I)} \in E^{(I)}$, the rejection function is given by

$$s(a) = \frac{q(a)}{\sup_{\ell \in \mathcal{A}} q(\ell)} = 1,$$

thus we reject with probability 0 any sample that is completable. The expected runtime is therefore the inverse of the probability of generating a completable sample, i.e., $\mathbb{P}(X^{(I)} \in E^{(I)})^{-1}$, times the cost to generate a sample from $\mathcal{L}(X^{(I)})$. □

Algorithm 6 PDC deterministic second half with uniform weight

Input:
Distributions $\mathcal{L}(X_1), \ldots, \mathcal{L}(X_n)$;
measurable function T, $t \in \text{range}(T)$;
index set $I \subset \{1, \ldots, n\}$.
Output:
A sample from $\mathcal{L}((X_1, \ldots, X_n)|T = t)$.
Assumptions:
(A1), (A2), (DSH);
$q(a) = q(b)$ for all $a, b \in E^{(I)}$.

Sample from $\mathcal{L}(X^{(I)})$, denote the observation by $x^{(I)}$.
if $x^{(I)} \in E^{(I)}$ **then**
  return $\sigma_I(x^{(I)}, y_I)$
else
  restart
end if

See Section 7.7 for an example involving sums of independent uniform random variables.

7.2. **Exponential Distribution.** When the X_i, $i = 1, \ldots, n$, are independent and exponentially distributed random variables with parameters $\lambda_i > 0$, the marginal density function of X_i is given by

$$f_{X_i}(x) = \lambda_i e^{-\lambda_i x}, \qquad x > 0.$$

When the event E is of the form $E = \{\sum_{i=1}^n X_i = k\}$, then $\mathbb{P}(E) = 0$, and there is no hard rejection sampling algorithm in general. We take $I = \{i\}$, then since the density f_{T_i} is bounded by λ_i, we apply Algorithm 5. The acceptance condition is thus

$$\left\{ \sum_{j \neq i} X_j \leq k \text{ and } U < e^{-\lambda_i y_I} \right\},$$

where $y_I = k - \sum_{j \neq i} x_j$. By Theorem 5.2, we have

$$\text{Time}_{\text{CDSH}}(\mathbf{X}') = O\left(\frac{\lambda_i}{\mathbb{P}(\sum_{j \neq i} X_j \leq k)} \right).$$

7.3. **Beta Distribution.** A continuous random variable X is said to have Beta(α, β) distribution, $\alpha > 0, \beta > 0$, if it has density

$$f_X(x) = c_{\alpha,\beta}\, x^{\alpha-1}(1-x)^{\beta-1}, \qquad 0 < x < 1,$$

where $c_{\alpha,\beta}$ is the normalization constant. When at least one of α, β is less than 1, the density $f_X(x)$ is not bounded. When both α and β are at least 1, then we have

$$\max_x f_X(x) = \frac{\alpha - 1}{\alpha + \beta - 2}, \qquad \alpha, \beta > 1.$$

Thus, if we consider $X_1, X_2, \ldots, X_n$ independent Beta(α_j, β_j), $j = 1, \ldots, n$, with $E = \{\sum_{j=1}^n X_j = k\}$, then as long as there exists an index i such that both α_i and β_i are greater

than 1, we can apply Algorithm 5. The acceptance condition is given by

$$\left\{ \sum_{j \neq i} X_i \in [k, k-1] \text{ and } U < \frac{c_{\alpha,\beta}(\alpha+\beta-2)}{\alpha-1} y_I^{\alpha-1}(1-y_I)^{\beta-1} \right\},$$

where $y_I = k - \sum_{j \neq i} x_i$.

7.4. **Small Ball Probabilities.** Suppose X_i, $i \geq 1$ are i.i.d. with distribution $\mathbb{P}(X_i = 1) = \mathbb{P}(X_i = -1) = \frac{1}{2}$. Let w_i denote real–valued weights with $|w_i| \geq 1$, $i \geq 1$. Define $T := \sum_{i=1}^n w_i X_i$. Then for some open set G, $\mathbb{P}(T \in G)$ is known as the small ball probability, see for example [42]. To obtain sample paths, the simplest approach is to apply hard rejection sampling. However, noting that X_i is actually a discrete uniform distribution over the set $\{-1, 1\}$, using Algorithm 6 we can apply PDC deterministic second half and select any index I, and sample until $\{X^{(I)} \in G^{(I)}\}$, where $G^{(I)} = \bigcup_{g \in G}(g + w_I) \cup (g - w_I)$.

When G is an open set of length $2r$, then it was shown in [27] that $P(T \in G)$ is at most 2^{-n} times the sum of the largest r binomial coefficients in n. Let us assume that the values w_i are all integer–valued, and $G = (-1, 1)$. Let $S_r(n)$ denote the sum of the largest r binomial coefficients in n. We have

$$\mathbb{P}\left(T \in (-1,1)\right) \leq 2^{-n} S_1(n) = 2^{-n} \binom{n}{\lfloor n/2 \rfloor}.$$

If we apply PDC deterministic second half, this becomes

$$\mathbb{P}\left(X^{(I)} \in (-w_I - 1, -w_I + 1) \cup (w_I - 1, w_I + 1)\right) \leq 2^{-(n-1)} S_2(n-1) = 2^{-(n-1)} \binom{n}{\lfloor n/2 \rfloor},$$

which saves at most an anticipated factor of 2.

Also, as was exploited in [4], if in addition there exist two distinct elements $w_j \neq w_\ell$, then we let $I = \{j, \ell\}$, and we have the range of $w_j X_j + w_\ell X_\ell$ is uniform over four distinct elements, say $\{v_1, v_2, v_3, v_4\}$; let $V_i = (v_i - 1, v_i + 1)$, then $V := \cup_{i=1}^4 V_i$ is an open set of length 8, whence

$$\mathbb{P}(T - w_j X_j - w_\ell X_\ell \in V) \leq 2^{-(n-2)} S_4(n-2).$$

One can keep going with this idea. If we let $I = \{j_1, j_2, \ldots\}$, then we must have that $\sum_\ell w_{j_\ell} X_{j_\ell}$ is uniform over distinct elements; i.e., each combination of ± 1 in X_{j_ℓ} must yield a distinct element for the sum. This is true, e.g., if $w_{j_1} = 1$, $w_{j_2} = 2$, $\ldots$, $w_{j_\ell} = 2^\ell$.

7.5. **Sampling from the surface of the n–sphere.** The following is *not* an example of PDC deterministic second half, but serves to illustrate the versatility of the PDC approach. Consider the distribution

$$\left(X_1, \ldots, X_n \,|\, X_1^2 + \ldots + X_n^2 = k\right),$$

where all random variables are continuous, which corresponds to some distribution on the surface of an n–sphere. It is known how to obtain the uniform distribution and certain other distributions over the surface of an n–sphere, see e.g. [20, 21] (see also [8] for a generalization to the ℓ_p^n-ball); however, if we change the form of the conditioning slightly (for example, replace X_1^2 with X_1 in the conditioning event), or place a particular demand on any of the marginal laws of $X_1, \ldots, X_n$, then these techniques do not generalize in a straightforward manner. PDC, on the other hand, is robust with respect to small changes.

We take $I = \{i\}$, for some $i \in \{1, \ldots, n\}$. Then we have

$$\mathcal{L}(T_I \,|\, X^{(I)} = x^{(I)}, E) = \mathcal{L}(X_i^2 \,|\, X_i^2 = t_I(x^{(I)})).$$

The rejection has the form:

$$\left\{ y_I \in \text{range}(X_i^2) \text{ and } U < \frac{\frac{1}{2\sqrt{t_I}} \left(f_{X_i}(\sqrt{t_I}) + f_{X_i}(-\sqrt{t_I}) \right)}{\sup_\ell f_{X_i^2}(\ell)} \right\}.$$

Note that we have calculated explicitly the transformation from the distribution $\mathcal{L}(T_I \,|\, X^{(I)})$ to the distribution $\mathcal{L}(X_I)$. This is *almost* PDC deterministic second half, see [17, 18], because there are actually two possible values for X_I to complete the sample, even though the distribution $\mathcal{L}(T_I \,|\, X^{(I)})$ is trivial. Once a sample is accepted, we simply choose an outcome, $\sqrt{t_I}$ or $-\sqrt{t_I}$, in proportion to its value determined by the density function f_{X_i}.

This example also illustrates why the deterministic second half condition (DSH) is a statement about the number of ways to complete a sample given $X^{(I)} \in E^{(I)}$, rather than a statement about the triviality of the distribution $\mathcal{L}(T_I \,|\, X^{(I)})$.

7.6. **Uniform Spacings.** Suppose we place m points uniformly distributed over the interval $[0, 1]$, call them $u_1, u_2, \ldots, u_m$. Let $u_{(1)}$, $u_{(2)}$, $\ldots$, $u_{(m)}$ denote the ordering of the points in ascending order. Then it is well–known, see for example [20, 30], that the marginal distributions are given by

$$u_{(i)} \sim \text{Beta}\,(i, m + 1 - i)\,, \qquad i = 1, \ldots, n.$$

The differences between consecutive points in the interval are in fact i.i.d. with (taking $u_{(0)} = 0, u_{(m+1)}) = 1$)

$$u_{(i)} - u_{(i-1)} \stackrel{D}{=} \frac{E_i}{\sum_{i=1}^{m+1} E_i}, \qquad i = 1, \ldots, m + 1,$$

where E_i are exponential distributions with parameter 1. In other words, to obtain a sample from $(E_1, E_2, \ldots, E_n \,|\, \sum_{i=1}^n E_i = 1)$, one can follow the steps above, which does not use PDC. Of course, as alluded to in Section 7.5, any deviation from this very specific form of distribution renders this approach effectively useless, whereas PDC can still be applied.

7.7. **Convex Polytope Sampling.** Suppose we wish to sample from a convex polytope $P \subset \mathbb{R}^n$ with vertices $\{v_1, \ldots, v_m\}$. Assuming none of the points v_i are degenerate, i.e., there do not exist any v_i such that $v_i \in \text{ConvexHull}\{v_1, \ldots, v_{i-1}, v_{i+1}, \ldots, v_m\}$, we can sample uniformly from P via Algorithm 7 (see [30]; see also [20]).

Algorithm 7 [30] Convex polytope sampling

Assumptions: P is a convex polytope with m vertices.
Generate $(U_1, U_2, \ldots, U_{m-1})$ i.i.d in the interval $[0, 1]$, denoted by $(u_1, \ldots, u_{m-1})$.
Sort the points and denote them as $u_{(1)}, \ldots, u_{(m-1)}$. Let $u_{(0)} = 0$ and $u_{(m)} = 1$.
Let $y_i = u_{(i)} - u_{(i-1)}$ for $i = 1, \ldots, m$.
return $\sum_{i=1}^m y_i \, v_i$.

The key aspect of this algorithm is that it has a time and memory requirement on the order of m, the number of vertices, which means that it is not efficient for polytopes with a large number of vertices compared to their dimension, which we now demonstrate.

The hypersimplex $H_{n,k}$ is defined as

$$H_{n,k} = \left\{ (x_1, \dots, x_n) \in [0,1]^n : \sum_{i=1}^n x_i = k \right\}.$$

We can obtain a random point inside the hypersimplex using Section 7.1. Each coordinate is uniformly distributed over the interval $[0,1]$, so by Theorem 7.1, the PDC deterministic second half algorithm is simply to sample $(u_2, \dots, u_n)$ from independent uniform distributions over $[0,1]$ until $u_1 := k - \sum_{i=2}^n u_i \in [0,1]$.

The Permutahedron P_n, see e.g., [48], is the convex hull of all $m = n!$ permutations of the coordinates of the point $(1, 2, \dots, n)$. It can be described as follows:

$$P_n = \left\{ (x_1, \dots, x_n) \in [1,n]^n : \sum_{i=1}^n x_i = \binom{n+1}{2}, R \right\},$$

where R denotes Rado's condition [52], which is

$$R = \{\text{for all } j \geq 1,\ x_{(n)} + \ldots + x_{(j)} \leq n + (n-1) + \ldots + j\}.$$

Let

$$Q_n = \left\{ (x_1, \dots, x_n) \in [1,n]^n : \sum_{i=1}^n x_i = \binom{n+1}{2} \right\}.$$

This is a scaled version of the hypersimplex, which can be sampled using Section 7.1 where each coordinate is uniformly distributed over the interval $[1,n]$. Its asymptotic volume is given in [12] as

$$\mathrm{Vol}(Q_n) \sim (n-1)^n \sqrt{\frac{6}{\pi n}}.$$

It is well–known that $\mathrm{Vol}(P_n) = n^{n-2}$, see for example [48, Proposition 2.4], i.e., the number of forests on n labeled vertices. Hence, the probability that a point in Q_n is also in P_n is given by

$$\mathbb{P}(R|Q_n) \sim \frac{n^{n-2}}{(n-1)^n \sqrt{\frac{6}{\pi n}}} \sim n^{-3/2}\, e \sqrt{\frac{\pi}{6}}.$$

Thus, to sample points from inside the permutahedron, one can first sample points from inside the scaled hypersimplex using Algorithm 6, and then apply hard rejection sampling with respect to Rado's condition R.

8. Applications to combinatorial classes

8.1. **Table Methods.** In any discussion of random sampling of combinatorial structures, invariably one is led to the recursive method of Nijenhuis and Wilf [44], which uses a table of values to calculate conditional probability distributions based on recursive properties of a corresponding combinatorial sequence. This method has several costs:

1. Computational cost to create the table
2. Storage cost to store the table
3. Computational cost to generate samples from the table.

If one is capable of handling items 1 and 2, then the table methods are typically the fastest known methods for random generation, since they provide fast lookups equivalent to unranking algorithms; a good survey of these and other similar algorithms is [55]; see also [16]. An extensive treatment of how to apply PDC in this case is given in [18], which offers a larger speedup at the expense of creating a partial table, and applies the recursive method to a smaller-sized set of objects. The advantage of our current approach is that it is table-free and simple to implement, while still offering an improvement over hard rejection sampling.

8.2. **Assemblies, Multisets, and Selections.** Let $w : \mathbb{N} \to \mathbb{R}$ denote a weighting function, and define the sequence $w_i := w(i)$, $i \geq 1$, and we interpret $a \cdot b$ as the usual dot product. For the remainder of this section, in order to keep the notation in [7], we define $\mathbf{Z} \equiv \mathbf{X}$ and let $T = \mathbf{Z} \cdot w$. In the examples that follow, the random vector $\mathbf{Z}$ is discrete, hence the event $\{T = t\}$ has strictly positive probability for each $t \in \text{range}(T)$.

A well-known example is that of the cycle decomposition of a permutation. Letting $C_i(n)$ denote the number of cycles of length i in a random permutation of n, it was shown in [54] that the joint distribution of all cycle lengths satisfies

$$\mathcal{L}(C_1(n), \ldots, C_n(n)) = \mathcal{L}\left((Z_1, \ldots, Z_n) \middle| \sum_{i=1}^{n} i Z_i = n\right),$$

where $Z_1, \ldots, Z_n$ are independent Poisson random variables with $\mathbb{E}Z_i = 1/i$, $i = 1, 2, \ldots, n$.

There are many combinatorial distributions that can be described in this fashion. In [7], a unified framework is presented which is applied to three main types of combinatorial structures. We recount some of the basic definitions, and refer the interested reader to [7] for a more inspiring exposition.

For $\mathbf{a} = (a_1, a_2, \ldots, a_n)$ a sequence of nonnegative integers, and weights $\mathbf{w} = (w_1, \ldots, w_n)$, let $N(n, \mathbf{a}, \mathbf{w})$ denote the number of combinatorial objects of weight n having a_i components of size w_i, $i = 1, \ldots, n$. Since each component has size w_i, the total contribution to the weight of the object by component i is $w_i\, a_i$, and summing over all i gives us the total weight $n = \sum_{i=1}^{n} w_i\, a_i$ of the object. Suppose $J \subset \{1, \ldots, n\}$. The examples of interest will have the following form:

$$N(n, \mathbf{a}, \mathbf{w}) = 1(\mathbf{w} \cdot \mathbf{a} = n) f(J, n) \prod_{i \in I} g_i(a_i),$$

for some functions f and g_i, $i \in J$, with

$$p(n) = \sum_{\mathbf{a} \in \mathbb{Z}_+^n} N(n, \mathbf{a}, \mathbf{w})$$

denoting the total number of objects of weight n. We now suppose that our combinatorial objects are chosen uniformly at random. Then the number of components of size i is a random variable, say with distribution C_i, $i = 1, \ldots, n$, and $C = (C_1, \ldots, C_n)$ is the joint distribution of *dependent* random component sizes that satisfies $C \cdot \mathbf{w} = n$. The distribution of C is given by

$$\mathbb{P}(C = \mathbf{a}) = 1(\mathbf{w} \cdot \mathbf{a} = n) \frac{f(J, n)}{p(n)} \prod_{i \in J} g_i(a_i). \tag{9}$$

For each $x > 0$, let independent random variables Z_i, $i \in J$, have distributions

$$\mathbb{P}(Z_i = k) = c_i(x)\, g_i(k)\, x^{w_i k}, \qquad i \in J, \tag{10}$$

where c_i, $i \in J$, are the normalization constants, given by

$$c_i = \left(\sum_{k \geq 0} g_i(k) x^{w_i k} \right)^{-1}.$$

Now we can state the following theorem.

Theorem 8.1 ([7])**.** *Assume $J \subset \{1, \dots, n\}$. Let $Z_J = (Z_i)_{i \in J}$ denote a vector of independent random variables with distributions given by Equation* (10)*. Let $C_J = (C_i)_{i \in J}$ denote the stochastic process of random component sizes with distribution given by Equation* (9)*. Then*

$$C_J =^d (Z_J | T = n).$$

Furthermore,

$$\mathbb{P}(T = n) = \frac{p(n)}{f(J,n)} x^n \prod_{i \in J} c_i(x). \tag{11}$$

Remark 8.1. The hard rejection sampling algorithm for such combinatorial classes has cost

$$\mathrm{Time}_{\mathrm{rej}}(Z_J) = O\left(\mathrm{Time}(Z_J) \, \frac{f(J,n)}{p(n)} x^{-n} \prod_{i \in J} c_i(x)^{-1} \right).$$

The form of the condition $\{T = n\}$ implies that the PDC deterministic second half algorithm has $|I| = 1$, i.e., I consists of a single index. By Theorem 5.1, we have for any $i \in J$,

$$\text{speedup} = \Omega\left(\max_k c_i(x) \, g_i(k) \, x^{w_i k} \right)^{-1}.$$

The optimal choice of i is one that minimizes this maximal probability. In fact, *any choice of i* will provide a speedup, which is an even more compelling reason to use PDC in this setting.

8.3. **The Boltzmann Sampler.** The Boltzmann sampler is a popular approach to random generation of combinatorial structures based on generating functions; see [23, 25, 31]. In terms of the structures introduced in Section 8.2, the Boltzmann sampler is equivalent to sampling from the distribution $\mathcal{L}(\mathbf{X})$, and accepting *all* samples, regardless of whether they lie in the set E. While this is not exact sampling, for the purposes of obtaining large scale characteristics, this approach has proven to be fruitful and reasonably accurate for many statistics of interest; see [22] and the references therein.

In terms of exact sampling, the typical recommended procedure is to perform hard rejection sampling on Boltzmann samples; see [22, Section 5.2]. Recently, several authors have noted various ways in which to obtain exact sampling using Boltzmann sampling; see [9, 36], based on an original approach used in [1]. The approach is equivalent to self-similar PDC (see [5] and also Remark 8.2), where one divides the sample space into two roughly equal-sized parts (in terms of randomness, not necessarily in terms of component-sizes), each a recursively defined copy of the original, samples each separately, and then pieces them together. The main caveat is the ability to compute the appropriate rejection probabilities, which are often in a form which can be approximated by local central limit theorems. In the case of Motzkin words in [1], the rejection probabilities were computable since they were written as a quotient of certain binomial coefficients.

While we envision the ultimate goal of random sampling of combinatorial structures to be self-similar PDC, the computing of rejection probabilities is a major impediment. Once such probabilities become accessible, PDC deterministic second half may not be as competitive. Until then, we stress that the main utility of PDC deterministic second half is the fact that *no local limit theorems or detailed knowledge of the sample space is required knowledge*, only the explicit calculation of the distribution and its maximum in (4); see Remark 8.1.

8.4. **Example: Integer Partitions.** A random unrestricted integer partition of (non–random) size n is described by $J = \{1, \ldots, n\}$, $w(i) = i$, $f(J, n) = 1, g_i(a_i) = 1$, $p(n)$ is the number of partitions of n, usually denoted by $p(n)$. Hence, for each $0 < x < 1$, we have normalization factors $c_i = (1 - x^i)$, so Equation (10) specializes to

$$\mathbb{P}(Z_i = k) = x^{ik}(1 - x^i), \qquad 0 < x < 1,\ i = 1, \ldots n,$$

i.e., Z_i is geometrically distributed with parameter $1 - x^i$. In this example, Z_i denotes the number of parts of size i in a random partition (of random size).

The probability that we generate a particular partition of n, with multiplicities $(c_1, \ldots, c_n)$, is given by

$$\mathbb{P}(Z_1 = c_1, \ldots, Z_n = c_n) = \prod_{i=1}^{n} \mathbb{P}(Z_i = c_i) = \prod_{i=1}^{n} x^{i\,c_i}(1 - x^i) = x^{\sum_{i=1}^{n} i\,c_i} \prod_{i=1}^{n} (1 - x^i).$$

Then, since each partition of n satisfies $\sum_{i=1}^{n} i\,c_i = n$, we have

$$\mathbb{P}(Z_1 = c_1, \ldots, Z_n = c_n) = x^n \prod_{i=1}^{n} (1 - x^i).$$

In other words, only the size of the partition determines its likelihood of being generated using this approach, and all partitions of the same size are equally likely to appear. Thus, either by summing over all partitions of n, or by Equation (11) directly, we have

$$\mathbb{P}(T = n) = p(n)\, x^n \prod_{i=1}^{n} (1 - x^i).$$

We would like to maximize this probability, and since this formula holds for all $0 < x < 1$, we can find an expression for x for which this expression is at its maximum. It was shown in [33] (see also [56]) that the choice $x = e^{-c/\sqrt{n}}$, $c = \pi/\sqrt{6}$, is particularly optimal, and produces

$$\mathbb{P}(T = n) \sim \frac{1}{\sqrt[4]{96}\, n^{3/4}}.$$

Thus, the hard rejection sampling algorithm has cost

$$\mathrm{Time}_{\mathrm{rej}}(\mathbf{X}') = O\left(\mathrm{Time}(\mathbf{X})\, n^{3/4}\right).$$

It was shown in [5] that for any selected index $i \in \{1, \ldots, n\}$, the PDC deterministic second half algorithm obtains a speedup to hard rejection sampling of size

$$\text{speedup} = \Omega\left(\max_{j \geq 0} \left(x^{ij}(1 - x^i)\right)\right)^{-1} = \Omega\left(1 - x^i\right)^{-1} = \Omega\left(\frac{\sqrt{n}}{i\,c}\right).$$

Thus, the optimal choice is $I = \{1\}$, and the total time for the PDC deterministic second half algorithm is

$$\mathrm{Time}_{\mathrm{DDSH}}(\mathbf{X}') = O\left(\mathrm{Time}(\mathbf{X})\, n^{1/4}\right).$$

Remark 8.2. It was shown in [5] that by using $I = \{1, 3, 5, \ldots\}$, i.e., all odd parts, one can apply a self–similar recursive PDC algorithm for which the total time of the algorithm is $O(\text{Time}(\mathbf{X}))$; this method is PDC but not PDC deterministic second half. The algorithm relies on a bijection and the fact that the rejection function can be computed to arbitrary accuracy efficiently using the analysis in [37, 39, 41, 51], and also the log-concavity of the partition function for all $n \geq 26$ due to [43]; see also [19]. Thus, for *unrestricted* integer partitions we prefer the self–similar algorithm, but in general such an advantageous structure may not be known.

When partitions are restricted to have distinct parts, then we have $J = \{1, \ldots, n\}$, $w(i) = i$, $f(J, n) = 1$, $g_i(a_i) = 1(a_i \leq 1)$, $p(n)$ is the number of partitions of n into distinct parts, usually denoted by $q(n)$. Then we have normalization constants $c_i = (1 + x^i)^{-1}$, so

$$\mathbb{P}(Z_i = k) = \frac{x^{ik}}{1 + x^{ik}}, \qquad k \in \{0, 1\},\ 0 < x < 1,\ i = 1, \ldots n,$$

i.e., Z_i is Bernoulli with parameter $\frac{x^i}{1+x^i}$. For any i, the speedup is thus

$$\text{speedup} = \Omega\left(\max_{k\in\{0,1\}} \frac{x^{i\,k}}{1 + x^{i\,k}}\right)^{-1} = \Omega\left(\frac{x}{1+x}\right)^{-1} = \Omega(1).$$

Thus, for partitions into distinct parts, PDC deterministic second half offers only a constant factor improvement.

8.5. **Example: Selections.** Integer partitions of size n into distinct parts is an example of a selection: each element $\{1, 2, \ldots, n\}$ is either in the partition or not in the partition. Selections in general allow m_i different types of a component with weight i. For integer partitions into distinct parts, this would be like assigning m_i colors to integer i, and allowing at most one component of size i of each color. Then we have for all $0 < x < 1$

$$\mathbb{P}(Z_i = k) = \binom{m_i}{k} \left(\frac{x^i}{1 + x^i}\right)^k \left(\frac{1}{1 + x^i}\right)^{m_i - k},$$

which is binomial. The PDC deterministic second half algorithm using $I = \{i\}$ has

$$\text{speedup} = \Omega\left(\max_i \mathbb{P}\left(Z_i = \frac{m_i x^i}{1 + x^i}\right)^{-1}\right) = \Omega\left(\sqrt{\frac{m_i\, x^i}{(1 + x^i)^2}}\right) = \Omega\left(\sqrt{\text{Var}(Z_i)}\right),$$

where the final two equalities follow since Z_i is a binomial random variable, whence the largest point probability is centered at its expectation, and is approximately $1/\sqrt{2\pi\text{Var}(Z_i)}$.

8.6. **Example: Multisets.** Unrestricted integer partitions of size n is an example of a multiset: each element $\{1, 2, \ldots, n\}$ can appear any number of times in the partition. Multisets in general allow m_i different types of a component with weight i, similar to selections. We have for all $0 < x < 1$

$$\mathbb{P}(Z_i = k) = \binom{m_i + k - 1}{k} (1 - x^i)^{m_i} x^{ik}, \qquad k = 0, 1, \ldots,$$

which is negative binomial. The mode of the negative binomial distribution is given by the mass at $\lfloor (m_i - 1)x^i/(1-x^i) \rfloor$. Similarly as with selections, we have

$$\text{speedup} = \Omega\left(\max_i \mathbb{P}\left(Z_i = \frac{m_i x^i}{1-x^i}\right)^{-1}\right) = \Omega\left(\sqrt{\frac{m_i\, x^i}{(1-x^i)^2}}\right) = \Omega\left(\sqrt{\text{Var}(Z_i)}\right).$$

8.7. **Assemblies.** Assemblies are described using Z_i as Poisson(λ_i), where $\lambda_i = \frac{m_i x^i}{i!}$, $i = 1, \ldots, n$, and where m_i denotes the number of components of size i, and $x > 0$. We have

$$\mathbb{P}(Z_1 = c_1, \ldots, Z_n = c_n) = \prod_{i=1}^{n} \frac{m_i^{c_i} x^{i\, c_i}}{i!^{c_i} c_i!} e^{-\lambda_i} = x^n e^{-\sum_{i=1}^n \lambda_i} \prod_{i=1}^{n} \frac{m_i^{c_i}}{i!^{c_i} c_i!}. \tag{12}$$

Here again since the random variables are Poisson, the local central limit theorem implies for λ_i large that we have

$$\text{speedup} = \Omega\left(\sqrt{\text{Var}(Z_i)}\right) = \Omega\left(\sqrt{\lambda_i}\right),$$

hence we should select the index i with the largest variance to obtain the largest speedup.

For set partitions, C_i denotes the number of blocks of size i in a random set partition of size n, $i = 1, \ldots, n$. In this case $m_i = 1$, and we have $\lambda_i = x^i/i!$, for $x > 0$. In this case the probability of the event $\{T = n\}$ is maximized by the value of x such that $x\, e^x = n$, for which $x = \log(n)$ serves as a reasonable approximation for large n.[2] It was shown in [5] that $I = [\log(n)]$ is a particularly good choice, since

$$\lambda_I \sim \frac{\log(n)^{\log(n)}}{(\log(n))!} \sim \frac{e^{\log(n)}}{\sqrt{2\pi \log(n)}} \sim \frac{n}{\sqrt{2\pi \log(n)}},$$

hence,

$$\text{speedup} = \Omega\left(\sqrt{2\pi\lambda_I}\right) = \Omega\left(\sqrt{n}/\sqrt[4]{\log(n)}\right).$$

It fact, it was shown in [47] that

$$\mathbb{P}(T = n) = O\left(\sqrt{n\, \log(n)}\right),$$

whence the total time for PDC deterministic second half is $O(\log^{5/4}(n))$. One can even improve upon this rejection rate using the recursive method, see [18].

8.8. **Plane partitions.** We now recall the third example of [5, Section 3.3.1], which improves upon an algorithm in [11] for random sampling of plane partitions. The first step of the algorithm is to sample from a rectangular grid of random variables $\{Z_{i,j}(x)\}_{1 \le i,j \le n}$, where $Z_{i,j}$ is geometrically distributed with parameter x^{i+j+1}, with $x \sim 1 - (2\xi(3)/n)^{1/3}$ chosen to maximize the probability of the event $E = \left\{\sum_{i,j \ge 1} (i+j+1) Z_{i,j} = n\right\}$.

[2] In [15], it is shown that

$$x \sim \log(n) - \log(\log(n)) + O\left(\frac{\log\log n}{\log n}\right).$$

Lemma 8.1. [11, Lemma 7] *Let $\boldsymbol{X} = \{Z_{i,j}(x)\}_{1\le i,j\le n}$, where $Z_{i,j}$ is geometrically distributed with parameter x^{i+j+1}, with $x \sim 1-(2\xi(3)/n)^{1/3}$; let $E = \left\{\sum_{i,j\ge 1}(i+j+1)Z_{i,j} = n\right\}$. Random sampling of $\mathcal{L}(\boldsymbol{X})$ can be performed in $O(n^{2/3})$ operations. Furthermore, the hard rejection sampling algorithm, i.e., the Boltzmann sampler with a hard rejection step, has probability $O(n^{2/3})$ of hitting the target. Therefore, sampling of $\mathcal{L}(\boldsymbol{X}')$ has total $O(n^{4/3})$ operations.*

To obtain a plane partition, they apply a bijection in [45], which only requires $O(n\log^3(n))$ operations, which makes the entire algorithm on the same order as the sampling of $\mathcal{L}(\mathbf{X}')$, namely, $O(n^{4/3})$. We surmise a self-similar PDC sampling algorithm for $\mathcal{L}(\mathbf{X}')$ exists which is $O(n^{2/3})$, however, for this particular application, this is not necessary. Consider the PDC with partition $A = \{Z_{i,j}\}_{i,j\ne 1}$, and $B = \{Z_{1,1}\}$. By Equation (8), the speedup over hard rejection sampling is given by

$$\text{speedup} = \Omega\left(\max_k \mathbb{P}(Z_{1,1}=k)\right)^{-1} = \Omega\left(\mathbb{P}(Z_{1,1}=0)\right)^{-1} = \Omega\left(\frac{1}{1-x}\right) = \Omega\left(\frac{n^{1/3}}{2\xi(3)}\right).$$

Thus, using PDC deterministic second half, sampling from $\mathcal{L}(\mathbf{X}')$ can be performed in $O(n)$ total operations, which makes the entire algorithm $O(n\log^3(n))$.

8.9. **Further restrictions on components.** In addition to the condition that $\{\sum_i w_i Z_i = n\}$, one can further condition, e.g., on the event $\{\sum_i Z_i = k\}$, which demands that the total number of components in the random structure with weight n is k; see, e.g., [7, Section 8]. For integer partitions, this is equivalent to sampling from partitions of size n with exactly k parts; or, set partitions of size n with exactly k blocks. In general, this condition takes the form $\{\sum_i u_i Z_i = k\}$ for some set of nonnegative coefficients $u = (u_i)_{i\ge 1}$. Then, for any $\theta > 0$, we let Z_i have distribution given by

$$\mathbb{P}_\theta(Z_i = c_i) = \frac{\theta^{u_i c_i}}{\mathbb{E}\theta^{u_i Z_i}}\mathbb{P}(Z_i = c_i), \qquad i = 1,\dots,n,$$

where we assume $\mathbb{E}\theta^{u_i Z_i} < \infty$. Then, by choosing a value of θ, we can effectively tilt the distribution.

For example, take $u_i = 1$ for all i and Z_i to be Poisson with parameter $\lambda_i = \frac{\theta x^i}{i}$, $i = 1,\dots,n$, and some $\theta > 0$. Then

$$\begin{aligned}\mathbb{P}(Z_1 = c_1,\dots,Z_n = c_n) &= \frac{x^n e^{-\theta\sum_{i=1}^n \lambda_i}}{n!}\frac{n!}{\theta(\theta+1)(\theta+2)\dots(\theta+n-1)}\prod_{i=1}^n \frac{\theta^{c_i}}{i^{c_i}c_i!} \\ &= \frac{\theta^k}{\theta(\theta+1)(\theta+2)\dots(\theta+n-1)}\frac{x^n e^{-\theta\sum_{i=1}^n \lambda_i}}{n!}n!\prod_{i=1}^n \frac{1}{i^{c_i}c_i!}.\end{aligned} \tag{13}$$

The joint distribution $\mathcal{L}(Z_1,\dots,Z_n)$ is the independent process approximation to the Ewens sampling formula, which is given similarly by $\mathcal{L}((Z_1,\dots,Z_n)|E)$, where $E = \{\sum_{i=1}^n iZ_i = n,\ \sum_{i=1}^n Z_i = k\}$, see [28]; see also [7, Section 8]. Letting $t_A := \sum_{i=3}^n i\,Z_i$ and $s_A := \sum_{i=3}^n Z_i$, there is a PDC deterministic second half algorithm for this family; namely, let $I = \{1,2\}$, then we have

$$E = \left\{\sum_{i=1}^n iZ_i = n,\ \ \sum_{i=1}^n Z_i = k\right\},$$

$$E^{(I)} = \left\{ (x_3, \ldots, x_n) : \sum_{i=3}^{n} i\, x_i \le n, \quad \sum_{i=3}^{n} x_i \le k \right\},$$

$$E_I = E_{I|X^{(I)}} = \left\{ (y, z) : y + 2z + \sum_{i=3}^{n} i\, Z_i = n, \quad y + z + \sum_{i=3}^{n} Z_i = k \right\},$$

$$\mathcal{L}\left(X^{(I)} \,\middle|\, E\right) = \mathcal{L}\left((Z_3, Z_4, \ldots, Z_n) \,\middle|\, \sum_{i=1}^{n} i\, Z_i = n,\ \sum_{i=1}^{n} Z_i = m \right),$$

$$\mathcal{L}\left(T_I \,\middle|\, X^{(I)} = x^{(I)}, E\right) = \mathcal{L}\left((Z_1, Z_2) \,\middle|\, Z_1 + 2Z_2 = n - t_A,\ Z_1 + Z_2 = m - s_A \right).$$

Thus, conditional on accepting a set of values for $X^{(I)} = (Z_3, \ldots, Z_n)$, the values of Z_1 and Z_2, say y_I and z_I, are uniquely determined, hence deterministic. The rejection condition is then given by

$$\left\{ X^{(I)} \in E^{(I)} \text{ and } U < \frac{\mathbb{P}(Z_1 = y_I)\mathbb{P}(Z_2 = z_I)}{\max_{j_1, j_2} \mathbb{P}(Z_1 = j_1)\mathbb{P}(Z_2 = j_2)} \right\},$$

and the speedup is

$$\text{speedup} = \Omega\left(\left(\max_{j_1, j_2} \mathbb{P}(Z_1 = j_1)\mathbb{P}(Z_2 = j_2) \right)^{-1} \right).$$

Again, we strongly note that no local central limit theorems are required, nor any qualitative or quantitative information pertaining to the Ewen's sampling formula or its approximation; the only requirement is the ability to efficiently compute the largest point probability of a given distribution.

There is also a two–parameter family which generalizes the Ewen's sampling formula given in [46]. In general, the same kind of PDC deterministic second half algorithm can be applied to any set of restrictions, as long as they satisfy the regularity conditions, and the choice of I uniquely determines a completion in the second stage of PDC.

8.10. **Logarithmic Combinatorial Structures.** While PDC deterministic second half is useful in full generality, and when limited information about the sample space is known, Remark 8.2 demonstrates that it is possible to substantially improve on the PDC deterministic second half approach for the sampling of $\mathcal{L}(\mathbf{X}')$ in certain cases. One particular class of examples is known as *logarithmic combinatorial structures*, see for example [3], which is the case

$$i\, \mathbb{E}X_i \to \kappa, \qquad i\, \mathbb{P}(X_i = 1) \to \kappa.$$

For assemblies, this is equivalent to $m_i \sim \kappa\, y^i (i-1)!$ as $i \to \infty$, for some $y > 0$ and $\kappa > 0$; for selections and multisets, this is equivalent to $m_i \sim \kappa\, y^i / i$ as $i \to \infty$, for some $y > 1$ and $\kappa > 0$.

One particular example, which demonstrates the utility of a more detailed understanding of the underlying structure, is for random sampling of permutations according to cycle structure. Let $C_i(n)$ denote the number of cycles of length i in a random permutation of n, and let X_i be Poisson with parameter $1/i$, $i = 1, 2, \ldots$. A classical result, see [35, 40], is that as n tends to infinity, we have

$$(C_1(n), C_2(n), \ldots) \xrightarrow{D} (Z_1, Z_2, \ldots).$$

The Feller coupling is a way to sample from $(C_1(n), C_2(n), \ldots)$ directly; see [29, 53]; see also [2, 3]. Start with 1 by itself in a cycle. With probability $1/n$, close off the cycle, and start a new cycle with 2 by itself; otherwise, add a uniform number $\{2, \ldots, n\}$ to the cycle with 1. Continue this process, closing off the cycle with probability $1/(n-1)$; otherwise, adding a number from the elements remaining uniformly at random to the currently open cycle. We have

$$C_i(n) = \#\{\text{exactly } i-1 \text{ times not closing a cycle followed by closing the cycle}\}.$$

According to our costing scheme, which ignores the sizes of the values, this algorithm is $O(n)$, which is asymptotically best possible[3].

This example serves to remind the reader that while PDC deterministic second half offers an improvement to hard rejection sampling, and is often just as practical given little knowledge of the underlying sample space, an improved algorithm can often be fashioned using more sophisticated arguments.

9. Acknowledgements

We gratefully acknowledge Richard Arratia, Tom Liggett, and Georg Menz for helpful suggestions, and Igor Pak for introducing us to the area of combinatorial polytopes and help with the literature. We also thank an anonymous referee for helpful comments which led to the improvement of Section 8.

References

[1] Laurent Alonso. Uniform generation of a Motzkin word. *Theoret. Comput. Sci.*, 134(2):529–536, 1994.
[2] Richard Arratia, A. D. Barbour, and Simon Tavaré. Poisson process approximations for the Ewens sampling formula. *Ann. Appl. Probab.*, 2(3):519–535, 1992.
[3] Richard Arratia, Andrew D Barbour, and Simon Tavaré. *Logarithmic combinatorial structures: a probabilistic approach*, volume 1. European Mathematical Society Zürich, 2003.
[4] Richard Arratia and Stephen DeSalvo. On the singularity of random bernoulli matrices—novel integer partitions and lower bound expansions. *Annals of Combinatorics*, 17(2):251–274, 2013.
[5] Richard Arratia and Stephen DeSalvo. Probabilistic divide-and-conquer: a new exact simulation method, with integer partitions as an example. *Combinatorics, Probability and Computing*, 25(3):324–351, May 2016.
[6] Richard Arratia and Simon Tavare. The cycle structure of random permutations. *The Annals of Probability*, pages 1567–1591, 1992.
[7] Richard Arratia and Simon Tavare. Independent process approximations for random combinatorial structures. *Adv. Math. 104 (1994), no. 1, 90-154*, 08 1994.
[8] Franck Barthe, Olivier Guédon, Shahar Mendelson, Assaf Naor, et al. A probabilistic approach to the geometry of the ℓ_p^n-ball. *The Annals of Probability*, 33(2):480–513, 2005.
[9] Olivier Bernardi and Omer Giménez. A linear algorithm for the random sampling from regular languages. *Algorithmica*, 62(1-2):130–145, 2012.
[10] Patrick Billingsley. *Convergence of probability measures*. John Wiley & Sons, Inc., New York-London-Sydney, 1968.
[11] Olivier Bodini, Éric Fusy, and Carine Pivoteau. Random sampling of plane partitions. *Combinatorics, Probability and Computing*, 19(02):201–226, 2010.
[12] L Carlitz, DC Kurtz, R Scoville, and OP Stackelberg. Asymptotic properties of eulerian numbers. *Probability Theory and Related Fields*, 23(1):47–54, 1972.
[13] George Casella and Roger L Berger. *Statistical inference*, volume 2. Duxbury Pacific Grove, CA, 2002.

[3]Note that while the entropy lower bound is $\log n! = \Omega(n \log(n))$, this assumes that the complexity of arithmetic manipulations of a number n grows like $O(\log(n))$, whereas we assume this cost is fixed at $O(1)$ for all n.

[14] Joseph T Chang and David Pollard. Conditioning as disintegration. *Statistica Neerlandica*, 51(3):287–317, 1997.
[15] Nicolaas Govert De Bruijn. *Asymptotic methods in analysis*, volume 4. Courier Dover Publications, 1970.
[16] Alain Denise and Paul Zimmermann. Uniform random generation of decomposable structures using floating-point arithmetic. *Theoretical Computer Science*, 218(2):233–248, 1999.
[17] Stephen DeSalvo. Exact, uniform sampling of latin squares and sudoku matrices. *arXiv preprint arXiv:1502.00235*, 2015.
[18] Stephen DeSalvo. Improvements to exact Boltzmann sampling using probabilistic divide-and-conquer and the recursive method. *arXiv preprint arXiv:1608.07922*, 2016.
[19] Stephen DeSalvo and Igor Pak. Log-concavity of the partition function. *The Ramanujan Journal*, 10 2013.
[20] Luc Devroye. Nonuniform random variate generation. *Handbooks in operations research and management science*, 13:83–121, 2006.
[21] Persi Diaconis and Laurent Saloff-Coste. Bounds for kac's master equation. *Communications in Mathematical Physics*, 209(3):729–755, 2000.
[22] Philippe Duchon. Random generation of combinatorial structures: Boltzmann samplers and beyond. In *Proceedings of the 2011 Winter Simulation Conference (WSC)*, pages 120–132. IEEE, 2011.
[23] Philippe Duchon, Philippe Flajolet, Guy Louchard, and Gilles Schaeffer. Random sampling from boltzmann principles. In *Automata, languages and programming*, pages 501–513. Springer, 2002.
[24] Philippe Duchon, Philippe Flajolet, Guy Louchard, and Gilles Schaeffer. Boltzmann samplers for the random generation of combinatorial structures. *Combinatorics, Probability and Computing*, 13(4-5):577–625, 2004.
[25] Philippe Duchon, Philippe Flajolet, Guy Louchard, and Gilles Schaeffer. Boltzmann samplers for the random generation of combinatorial structures. *Combin. Probab. Comput.*, 13(4-5):577–625, 2004.
[26] Rick Durrett. *Probability: theory and examples*. Cambridge university press, 2010.
[27] P. Erd?s and M Szalay. On the statistical theory of partitions, 1981.
[28] W. J. Ewens. The sampling theory of selectively neutral alleles. *Theoret. Population Biology*, 3:87–112; erratum, ibid. 3 (1972), 240; erratum, ibid. 3 (1972), 376, 1972.
[29] W. Feller. The fundamental limit theorems in probability. *Bull. Amer. Math. Soc.*, 51:800–832, 1945.
[30] William Feller. An introduction to probability and its applications, vol. ii. *John Wiley, New York*, 1971.
[31] Philippe Flajolet and Robert Sedgewick. *Analytic combinatorics*. Cambridge University Press, Cambridge, 2009.
[32] Gerald B Folland. *Real analysis: modern techniques and their applications*. John Wiley & Sons, 2013.
[33] Bert Fristedt. The structure of random partitions of large integers. *Transactions of the American Mathematical Society*, 337(2):703–735, 1993.
[34] Stuart Geman and Donald Geman. Stochastic relaxation, gibbs distributions, and the bayesian restoration of images. *IEEE Transactions on pattern analysis and machine intelligence*, 6(6):721–741, 1984.
[35] VL Goncharov. Some facts from combinatorics. *Izvestia Akad. Nauk. SSSR, Ser. Mat*, 8:3–48, 1944.
[36] Dominique Gouyou-Beauchamps and Cyril Nicaud. Random generation using binomial approximations. In *21st International Meeting on Probabilistic, Combinatorial, and Asymptotic Methods in the Analysis of Algorithms (AofA'10)*, Discrete Math. Theor. Comput. Sci. Proc., AM, pages 359–371. Assoc. Discrete Math. Theor. Comput. Sci., Nancy, 2010.
[37] Godfrey Harold Hardy and Srinivasa Ramanujan. Asymptotic formulaæ in combinatory analysis. *Proceedings of the London Mathematical Society*, 2(1):75–115, 1918.
[38] Mark L. Huber. *Perfect Simulation*. Chapman & Hall/CRC Monographs on Statistics & Applied Probability. Taylor & Francis, 2015.
[39] Fredrik Johansson. Efficient implementation of the hardy–ramanujan–rademacher formula. *LMS Journal of Computation and Mathematics*, 15:341–359, 2012.
[40] VF Kolchin. A problem of the allocation of particles in cells and cycles of random permutations. *Theory of Probability & Its Applications*, 16(1):74–90, 1971.
[41] Derrick Henry Lehmer. On the remainders and convergence of the series for the partition function. *Transactions of the American Mathematical Society*, 46(3):362–373, 1939.

[42] Hoi H Nguyen and Van H Vu. Small ball probability, inverse theorems, and applications. In *Erdős Centennial*, pages 409–463. Springer, 2013.
[43] Jean-Louis Nicolas. Sur les entiers N pour lesquels il y a beaucoup de groupes abéliens d'ordre N. *Ann. Inst. Fourier (Grenoble)*, 28(4):1–16, ix, 1978.
[44] Albert Nijenhuis and Herbert S. Wilf. A method and two algorithms on the theory of partitions. *J. Combinatorial Theory Ser. A*, 18:219–222, 1975.
[45] Igor Pak. Hook length formula and geometric combinatorics. *Sém. Lothar. Combin.*, 46:Art. B46f, 13 pp. (electronic), 2001/02.
[46] Jim Pitman. The two-parameter generalization of ewens' random partition structure. Technical report, Technical Report 345, Dept. Statistics, UC Berkeley, 1992.
[47] Boris Pittel. Random set partitions: asymptotics of subset counts. *journal of combinatorial theory, Series A*, 79(2):326–359, 1997.
[48] Alexander Postnikov. Permutahedra, associahedra, and beyond. *arXiv preprint math.CO/0507163*, 2005.
[49] James Gary Propp and David Bruce Wilson. Exact sampling with coupled markov chains and applications to statistical mechanics. *Random structures and Algorithms*, 9(1-2):223–252, 1996.
[50] Michael A Proschan and Brett Presnell. Expect the unexpected from conditional expectation. *The American Statistician*, 52(3):248–252, 1998.
[51] Hans Rademacher. On the partition function $p(n)$. *Proceedings of the London Mathematical Society*, 2(1):241–254, 1938.
[52] R. Rado. An inequality. *J. London Math. Soc.*, 27:1–6, 1952.
[53] Alfréd Rényi. On the outliers of a series of observations. *Magyar Tud. Akad. Mat. Fiz. Oszt. Közl.*, 12:105–121, 1962.
[54] L. A. Shepp and S. P. Lloyd. Ordered cycle lengths in a random permutation. *Trans. Amer. Math. Soc.*, 121:340–357, 1966.
[55] Dennis Stanton and Dennis White. *Constructive combinatorics*. Springer-Verlag New York, Inc., 1986.
[56] Harold N. V. Temperley. Statistical mechanics and the partition of numbers ii. the form of crystal surfaces. In *Mathematical Proceedings of the Cambridge Philosophical Society*, volume 48, pages 683–697. Cambridge Univ Press, 1952.
[57] Tue Tjur. *Conditional probability distributions*, volume 2. Institute of Mathematical Statistics, University of Copenhagen, 1974.
[58] John Von Neumann. Various techniques used in connection with random digits. *Applied Math Series*, 12(36-38):1, 1951.
[59] BB Winter. An alternate development of conditioning. *Statistica Neerlandica*, 33(4):197–212, 1979.

UCLA Department of Mathematics, 520 Portola Plaza, Los Angeles, CA, 90095
E-mail address: stephendesalvo@math.ucla.edu